Soft tensegrity robots

John Rieffel[1,*] and Jean-Baptiste Mouret[2,3,4,*]

[1]Union College, Schenectady, NY 12308, USA
[2] Inria Nancy Grand - Est, Villers-lès-Nancy, F-54600, France
[3]CNRS, Loria, UMR 7503, Vandœuvre-lès-Nancy, F-54500, France
[4] Universit de Lorraine, Loria, UMR 7503, Vandœuvre-lès-Nancy, F-54500, France
[*]J.R. and J.-B. M. contributed equally to this work

Preprint – February 13, 2017

arXiv:1702.03258v1 [cs.RO] 10 Feb 2017

Living organisms intertwine soft (e.g., muscle) and hard (e.g., bones) materials, giving them an intrinsic flexibility and resiliency often lacking in conventional rigid robots. The emerging field of soft robotics seeks to harness these same properties in order to create resilient machines. The nature of soft materials, however, presents considerable challenges to aspects of design, construction, and control – and up until now, the vast majority of gaits for soft robots have been hand-designed through empirical trial-and-error. This manuscript describes an easy-to-assemble tensegrity-based soft robot capable of highly dynamic locomotive gaits and demonstrating structural and behavioral resilience in the face of physical damage. Enabling this is the use of a machine learning algorithm able to discover novel gaits with a minimal number of physical trials. These results lend further credence to soft-robotic approaches that seek to harness the interaction of complex material dynamics in order to generate a wealth of dynamical behaviors.

UNLIKE machines, animals exhibit a tremendous amount of resilience, due in part to their intertwining of soft tissues and rigid skeletons. In nature, this suppleness leads to several compelling behaviors which exploit the dynamics of soft systems. Octopi, for example, are able to adaptively shape their limbs with "joints" in order to perform efficient grasping[1]. Jellyfish exploit their inherent elasticity in order to passively recover energy during swimming[2]. *Manduca sexta* caterpillars have a mid-gut which acts like a "visceral-locomotory piston" – sliding forward ahead of the surrounding soft tissues, shifting the animal's center of mass forward well before any visible exterior change[3].

Taking inspiration from the natural world, the field of soft robotics seeks to address some of the constraints of conventional rigid robots through the use of compliant, flexible, and elastic materials[7,8]. Trimmer *et al.*, for instance, construct soft robots from silicone rubber, using shape-memory alloy (SMA) micro-coil actuation, which can slowly crawl in controlled fashion[9] or roll in an uncontrolled ballistic fashion[10]. Similarly, research by Whitesides *et al.* uses pneumatic inflation to produce slow, dynamically stable crawling motions[11] as well as fast, but less controlled tentacle-like grippers[12], combustion-driven jumpers[13] and a self-contained microfluidic "octobot"[8].

Despite their advantages, soft-material robots are difficult to control by conventional means[7,11]. They are by their very nature high dimensional dynamic systems with an essentially infinite number of degrees of freedom. The elasticity and deformability which provide their appeal come at the cost of resonances and tight dynamic coupling between components[9], properties which are often avoided, or at least suppressed, in conventional engineering approaches to robotic design. This complexity precludes the use of many of the traditional kinematic and inverse-dynamics approaches to robotic control[14].

As a result, up until now, the locomotive gaits of most soft robots have been developed by hand through empirical trial-and-error[11]. This process can be both challenging and time consuming, particularly when seeking to fully exploit the dynamical complexity of soft mechanisms. Importantly, this manual process also prevents these robots from adapting their control strategy when the context changes, for instance when they encounter an unexpected type of terrain, or when they are physically damaged.

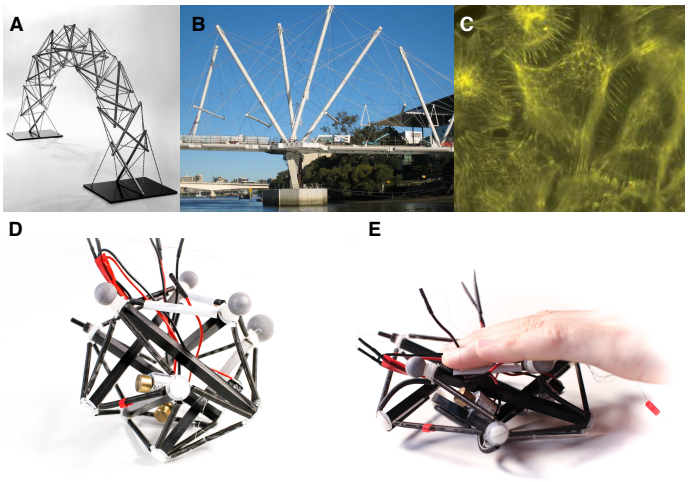


Fig. 1. Concept of our soft tensegrity robot. Tensegrity structures are combinations of rigid elements (struts) joined at their endpoints by tensile elements (spring or cables) that are kept stable by the interplay of pre-stress forces. **A**. The first tensegrity structures appeared in art, with the sculptures of Kenneth Snelson[4,5]. **B.** They have been subsequently used in architecture, for instance for the Kurilpa bridge (Brisbane, Australia). **C**. More recently, tensegrity has been found to be a good model of the mechanotransduction of living cells[6]. **D.** Our tensegrity robot is based on carbon struts and springs. It is is actuated by 3 vibrators (glued to 3 of the struts) whose frequency is automatically tuned by a trial-and-error learning algorithm (Methods). **E.** Thanks to the tensegrity structure and to the compliance of the springs, our robot will keep its integrity when deformed and spring back into its initial form. A video is available in supplementary materials (Video S1).

In this work, we introduce a new class of soft robot based upon a tensegrity structure driven by vibration. Like many other soft robots, this tensegrity robot is resilient, and can resist damage when perturbed or crushed. Unlike other soft robots, however, this particular modular tensegrity robot is easy to build, easy to control, and thanks to a data-efficient reinforcement learning algorithm[15], it can autonomously discover how to move, and quickly relearn and adapt its behavior when damaged.

Results

Tensegrities are relatively simple mechanical systems, consisting of a number of rigid elements (struts) joined at their endpoints by tensile elements (cables or springs), and kept stable through a synergistic interplay of pre-stress forces (Fig. 1 A-C). Beyond engineering, properties of tensegrity has been demonstrated at all scales of the natural world, ranging from the tendinous network of the human hand[16] to the mechanotransduction of living cells[6]. At every size, tensegrity structures exhibit two interesting features[4,5]: they have an impressive strength-to-weight ratio, and

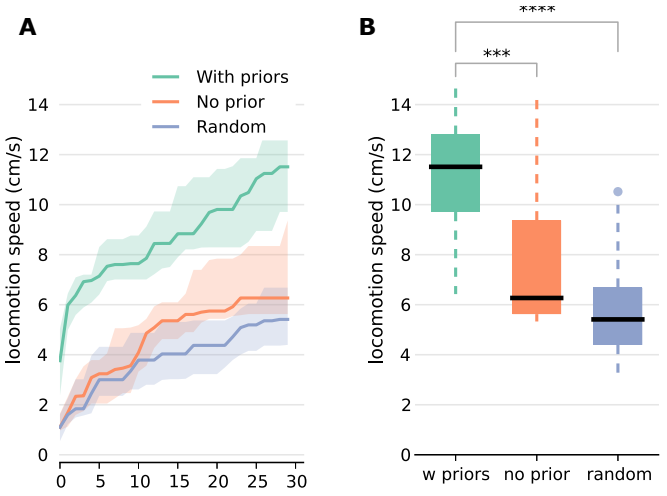


Fig. 2. Evaluation of the learning algorithm. A. Locomotion speed after each of the 30 trials. The light zones represent the 25th and 75th percentiles. **B. Locomotion speed after 30 trials.** The central mark is the median, the edges of the box are the 25th and 75th percentiles, the whiskers extend to the most extreme data points not considered outliers, and outliers are plotted individually. Each condition is tested with 20 independent runs of the algorithms.

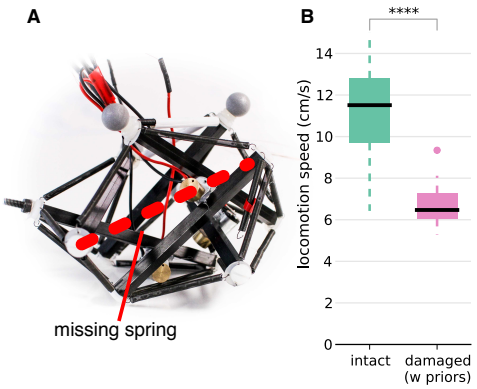


Fig. 3. Experiments with a damaged robot. A. Damaged robot. A spring is disconnected from the robot of Fig. 1. **B. Locomotion speed after 30 trials**. The central mark is the median, the edges of the box are the 25th and 75th percentiles, the whiskers extend to the most extreme data points not considered outliers, and outliers are plotted individually. Each condition is tested with 20 independent runs of the algorithms.

they are structurally robust and stable in the face of deformation. Moreover, unlike many other soft robots, tensegrity structures are inherently modular (consisting of only struts and springs) and are therefore relatively easy to construct. They are simple enough to be featured in books for children activities[17], while complex enough to serve as the basis for the next generation of NASA's planetary rovers[18].

Unfortunately, the control methods used in traditional robotics cannot be applied to tensegrity robots: the dynamical complexity and the high number of degrees of freedom make them very difficult to accurately model, simulate, and control[18]. The most common strategy is to slowly change the lengths of the struts and/or cables, which results in large-scale, quasi-static (rather than dynamic) structural deformations, which, in turn, make the robot move[18,19]. As they assume that the structure is stiff, such control strategies are not suitable for soft tensegrity robots. In addition, they lead to slow locomotion speeds.

More recently, researchers have begun investigating more dynamical methods of tensegrity robot control. Bliss *et al.* have used central pattern generators (CPGs) to produce resonance entrainment of simulated non-mobile tensegrity structures [20]. Mirletz *et al.* have used CPGs to produce goal-directed behavior in simulated tensegrity-spine-based robots [21]. These efforts, however valuable, were all produced in simulated environments, and have not yet been successfully transferred into real-world robots. As Mirletz *et al* point out [20], the dynamic behavior of tensegrities is highly dependent upon the substrate they interact with – this means that results developed in simulated environments cannot necessarily be simply transferred to real robots (in Evolutionary Robotics, this is known as the "Reality Gap" [22,23]).

Here we explore the hypothesis that the inherent resonance and dynamical complexity of real-world soft tensegrity robots can be beneficially harnessed (rather than suppressed), and that, if properly excited[24], it can resonate so that the robot performs step-like patterns that enable it to "walk". To test this hypothesis and demonstrate the potential of soft tensegrity robots, we designed a pocked-size, soft tensegrity robot whose parameters were tuned to maximize resonance, and whose goal is to locomote as fast as possible across flat terrain. To find the right resonances, we equipped the robot with a data-efficient trial-and-error algorithm, which also allows it to adapt when needed.

Our soft tensegrity robot (Fig. 1D-E) is based upon a canonical six-bar tensegrity shape consisting of equal length composite struts connected via 24 identical helical springs, with four springs emanating from each strut end. Unlike most tensegrity structures, which seek to maximize stiffness[24], the spring constants of our robot were chosen with the goal of producing suitably small natural frequencies of the structure, with corresponding large displacements – in other words, to maximize suppleness. This allows the pocket sized robot to maintain its structural shape under normal operation, and yet be easily compressed flat in one's hand. A variable speed motor coupled to offset masses was then attached to three of the struts in order to excite the natural frequencies of the structure.

Like many robots, the tensegrity robot needs to use different gaits in order to achieve lomocotion, depending on terrain. In our case, these gaits are determined by the speeds of the three vibratory motors. As the exact properties of the terrain are seldom known *a priori*, and because hand-designing gaits is time consuming (not to mention impossible when the robot is in remote or hazardous environments) this robot finds effective motor frequencies by using a trial-and-error learning algorithm whose goal is to maximize the locomotion speed.

Earlier work of ours [25,26] used interactive trial-and-error as well as automated hill climbing techniques to find optimal gaits for a tensegrity robot. These gaits, could in turn, be incorporated into a simple state machine for directional control. However, these techniques required hundreds of physical trials that were time consuming and produced significant wear on the physical robot. More importantly, the interactive procedure required an human in the loop, whereas we envision robots that can adapt autonomously to new situations (e.g. a damage or a new terrain).

Here, as a substantial improvement upon these earlier time-intensive intensive methods, we employ a Bayesian optimization algorithm[15,27,28], which is a mathematical optimizer designed to find the maximum of a performance function with as few trials as possible.

Conceptually, Bayesian optimization fits a probabilistic model (here a Gaussian process[29], see Methods) that maps motor speeds to locomotion speed. Because the model is probabilistic, the algorithm can not only predict which motor speeds are the most likely to be good, but also associate it to a confidence level. Bayesian optimization exploits this model to select the next trial by balancing exploitation – selecting motor speeds that are likely to make the robot move faster – and exploration – trying combinations of motor speeds that have not been tried so far (Methods). As an additional benefit, this algorithm can take into account that observations are by nature uncertain.

The Bayesian optimization algorithm usually starts with a con-

stant prior for the expected observation (e.g., the expected speed is 10 cm/s) and a few randomly chosen trials to initialize the model. For this robot, however, common sense, along with preliminary modeling, suggests that speeds near the motor maximums are more likely to produce successful gaits, and that near-zero motor speeds are not expected to make the robot move. This insight was substantiated in preliminary experiments: many effective gaits were produced by high motor speeds, both forward and backward. Therefore, to speed up learning, we use a non-linear prior model as follows: (1) if the three motor speeds are close to 0, then we should expect a locomotion speed close to 0 and (2) if all the motors are close to full speed (in any direction), then we should expect the maximum locomotion speed (Methods and Fig. 5D). Thanks to this prior, the Bayesian optimization algorithm does not need any random sample points to seed the prior and will instead start with promising solutions. In spite of this prior, learning is needed because many combinations of motors at full speeds make the robot tumble or rotate on itself, resulting in low performance; in addition, subtle changes to motor speeds can have dramatic effects upon the resulting robot gait.

We first evaluate the effectiveness of the learning algorithm (Fig. 2). The performance function is the locomotion speed, measured over 3 seconds, in any direction. If the robot turns too much, that is if the yaw exceeds a threshold, the evaluation is stopped (Methods). The covered distance is measured with an external motion capture system (Methods), although similar measurements can be obtained with an onboard visual odometry system[15,30]. We compare three algorithms: random search, Bayesian optimization without prior (using 10 random points to initialize the algorithm), and Bayesian optimization with prior. Each algorithm is allowed to test 30 different motor combinations (resulting in 90 seconds of learning for each experiment) and is independently run 20 times to gather statistics. The results show that the best locomotion speeds are obtained with the prior-based Bayesian optimization ($11.5cm/s$, 5$^{\text{th}}$ and 95$^{\text{th}}$ percentiles $[8.1, 13.7]$), followed by the prior-free Bayesian optimization ($6.3cm/s[5.5, 12.4]$). The worst results are obtained with the random search ($5.4cm/s[3.5, 9.9]$). Overall, these experiments demonstrate that the prior-based Bayesian optimization is an effective way to automatically discover a gait in only 30 trials with this robot. Videos of the gaits are available as supplementary material (Video S1).

We then investigate our hypothesis that the interplay between a flexible tensegrity structure and vibration is the key for effective locomotion. To do so, we designed a rigid replica of our robot that does not contain any springs: the carbon fiber struts are held in place with small carbon fiber rods (Fig. 4A). All the dimensions, strut positions, and motor positions are the same as for the tensegrity version (Fig. 1D-E). We used the same learning algorithm as for the tensegrity robot and replicated the learning experiment 20 times. The results (Fig. 4B) show that that the rigid replica is about 1.5 times slower than the tensegrity robot ($7.1cm/s[5.6, 9.3]$ vs $11.5cm/s[8.1, 13.7]$), which suggests that the flexibility of the tensegrity structure plays a critical role in its effective locomotion. In addition, we measured the amplitude of movement along the vertical axis for the end of 4 struts, both with the soft tensegrity robot and the rigid replica; we repeated this measure with 50 random gaits in both cases. These measurements (Fig. 4C) show that the markers move at least twice more when the structure is flexible ($2.3[1.5, 4.8]$ cm vs $0.99[0.61, 2.1]$ cm), which demonstrates that the structure amplifies the movements induced by the vibrators.

In addition of being deformable, tensegrity structures often maintain most of their shape when a link (a spring or a strut) is missing, leading to relatively smooth failure modes. We evaluate the ability of our robot to operate after such damage by removing a spring (Fig. 3A). As the shape of the robot is changed, we relaunch the learning algorithms. The experiments reveal that successful, straight gaits can be found in 30 trials, although they

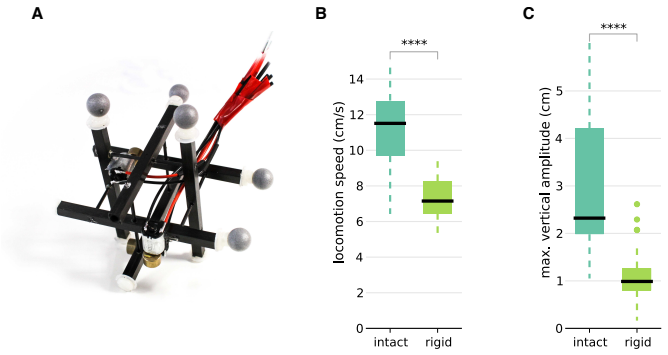


Fig. 4. Experiments with a rigid robot. A. Rigid replica of our soft tensegrity robot. This robot is identical to the robot of Fig. 1, except that it contains no spring: the carbon fiber struts are held in place with small carbon fiber rods. All the dimensions, strut positions, and motor positions are the same as for the tensegrity version. **B. Locomotion speed after 30 trials for the intact (Fig. 1) and the rigid robot (A).** Each condition is tested with 20 independent runs of the algorithms. **C. Maximum amplitude of the markers for random gaits**. In each case, we captured the vertical position of the 4 markers for 50 random gaits of 3 seconds. We report the maximum height minus the minimum height (over the 4 markers). For the box plots, the central mark is the median, the edges of the box are the 25$^{\text{th}}$ and 75$^{\text{th}}$ percentiles, the whiskers extend to the most extreme data points not considered outliers, and outliers are plotted individually.

are significantly lower-performing than gaits obtained with the intact robot ($11.5cm/s[8.1, 13.7]$ versus $6.5cm/s[5.6, 8.2]$ Fig. 3B).

During all the reported experiments, we evaluated $20 \times 30 \times 3 = 1800$ different gaits on the intact robot, $20 \times 30 = 600$ gaits on the rigid robot (20 replicates, 30 trials for each replicate, and 3 treatments), and $20 \times 30 = 600$ gaits on the damaged robot. We can use these points to draw a picture of the search space that does not depend on the learning algorithm (Fig. 5). Since the search space is too high-dimensional to be easily visualized (3 dimensions + performance, resulting in a 4D plot), we compute performance profiles[31,32]: for each combination of 2 motor speeds $\{v_1, v_2\}$, we report the best performance measured regardless of the speed of the third motor (Methods). The performance profiles (Fig. 5A) for the intact robot reveal that there are two high-performing regions, roughly positioned around $\{-100\%, 100\%, -100\%\}$ and $\{-100\%, -100\%, 100\%\}$ and that the first region ($\{-100\%, 100\%, -100\%\}$) is where most high-performing solutions can be found. This finding is consistent with the prior given to the learning algorithm (Fig. 5D), which models that the best performance should be obtained with a combination of -100% and $+100\%$ values. It should be emphasized that the best gaits do not correspond to the most extreme values for the motor speeds: the most reliable optima is around $\{-90\%, 100\%, -90\%\}$, mostly because too extreme values tend to make the robot tumble. The best solutions for the rigid robots are also found in the corners, that is, for combinations of $+100\%$ and -100% motor speeds, but the measurements suggest that the optimum might be different from the one obtained with the intact robot (more data would be needed to conclude). The data for the damaged robot show more clearly that the best solutions are around $\{-100\%, -100\%, 100\%\}$, which corresponds to the second optimum found for the intact robot (the lowest performing one).

The performance profiles thus demonstrate that the prior knowledge given to the learning algorithm is consistent with the three different robots (intact, rigid, and damaged), which suggests that it might be helpful in other situations (e.g., different damage conditions). They also demonstrate that gaits that work the best on the intact robot do not work on the damaged robot (Fig. 5 A versus C, second column): this shows that the learning algorithm is needed to adapt the gait if the robot is damaged.

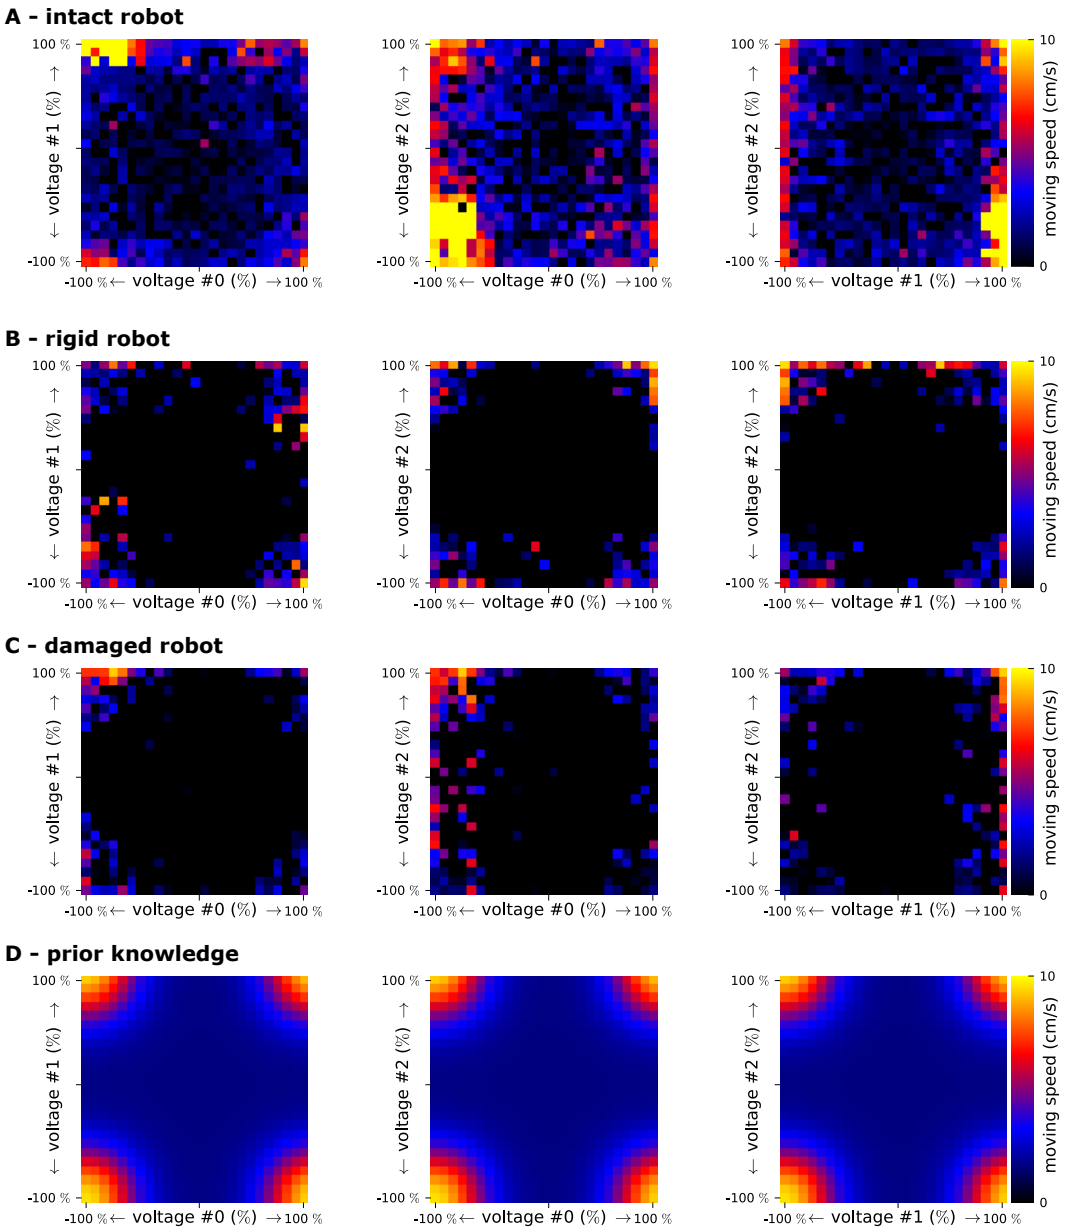


Fig. 5. Performance profiles for all the conditions. These performance profiles show the performance potential of each combination of 2 motor speeds (the third motor is considered as a "free variable"). Three plots are required to get a comprehensive picture of the performance space: v_1 vs v_2, v_1 vs v_3, and v_2 vs v_3. **A. Intact robot (Fig. 1D)**. The profiles are computed with 1800 policy evaluations (20 replicates $\times$ 30 trials $\times$ 3 sets of experiments – with prior, without prior, random search). **B. Rigid robot (Fig 4A).** The profiles are computed with 600 policy evaluations (30 trials $\times$ 20 replicates). **C. Damaged robot (Fig. 3)**. The profiles are computed with 600 policy evaluations (30 trials $\times$ 20 replicates). **D. Prior knowledge.** Prior knowledge used to guide the learning algorithm (Methods).

Discussion

Soft tensegrity robots are highly resilient, easy to assemble with the current technology, and made with inexpensive materials. Thanks to the learning algorithm, our prototype can achieve locomotion speeds of more than 10 cm/s (more 1 body length per second) and learn new gaits in less than 30 trials, which allows it to adapt to damage or new situations. To our knowledge, this makes it one of the fastest soft robot.

Our soft tensegrity robots achieve their speed because they harness the flexibility and the resonance of tensegrity structures, instead of trying to suppress it as it is done in architecture[5] or in other tensegrity robots[18,19]. Harnessing flexibility and resonance opens new research avenues for future tensegrity structures, in particular when mechanical design can be coupled with machine learning algorithms that automatically identify how to control the resonances.

As demonstrated here, vibration is an effective way to exploit structural flexibility; similar observations have been made with other, more rigid vibrating robots like the Kilobots[33]. A direct benefit is that it makes it easy to power soft tensegrity robots with an embedded battery, by contrast with the many fluid-actuated soft robots[7,11], which need innovative ways to store energy[8]. Nevertheless, soft tensegrity robots could excite their structure by other means; for instance, a flywheel that is rapidly decelerated could help the robot to achieve fast movements[34], or high-amplitude, low-frequency oscillations could be generated by moving a pendulum inside the structure[35].

Putting all these attractive features altogether, soft tensegrity robots combine simplicity, flexibility, performance, and resiliency, which makes this new class of robots one of the most promising building block for future soft robots.

Overall, soft tensegrity robots move thanks to the complex interactions between the actuators (vibrators), the structure (springs and struts), and the environment (the ground). As a consequence, they do not have a centralized control-policy like conventional robots; instead, their abilities are distributed throughout the robot. This kind of emergent behavior is central in the embodied intelligence theory[36], which suggests that we will achieve better and more life-like robots if we encourage such deep interactions between the body and the "mind" – here, the controller. However, as demonstrated in the present work, trial-and-error learning algorithms might be the only viable approaches to discover and adapt these emergent behaviors.

Material and Methods

Robot

The tensegrity used is defined by six equal length composite struts which are connected to each other via 24 identical helical springs, with four springs emanating from each strut end. This follows the geometry described as TR-6 by Skelton[5]. Few actual machining operations are required to produce the tensegrity. The six 9.4 cm long composite struts are cut from 6.35 mm square graphite composite tubes (Goodwinds). The three 12mm vibrational motors (Precision Microdrives Model 312-107) were mounted to the flat outer surface of the struts using hot melt adhesive. Both ends of each strut were then tapped for 10-24 nylon screws fitted with nylon washers. The hooked ends of the helical springs (Century Spring Stock No. 5368) were attached directly through the nylon washers. The motors were connected via thin gauge magnet wire to Serial Motor Controllers (Pololu Qik 2s9v1 Dual Serial Motor Controller) connected in turn to a USB Serial Adapter (SparkFun FTDI Basic Breakout board)

Control policy

Each policy is defined by three PWM values that determine the input voltage of the 3 vibrating motors ($\chi = [v_1, v_2, v_3]$), which can take values between 0 (full speed, backward) and 1 (full speed, forward); 0.5 corresponds to a speed of 0, that is, to no movement.

Performance function

Each controller is tested for 3 seconds, then the Euclidean distance between the starting point and the end point is recorded. The performance function is the distance (in cm/s) divided by 3. If during the 3 second evaluation period the yaw of the robot exceeds 1 radian, the evaluation is stopped and the recorded distance is the distance between the starting point and the point reached by the robot when it exceeded the yaw limit.

The policies are evaluated externally with a motion tracking system (Optitrack Prime 13 / 8 cameras), but the same measurements can be obtained with an embedded camera connected to a visual odometry system[15,30].

Profile plots

We use the profile plots to depict the search space and the prior used by the learning algorithm (Fig. 5). For each pair of dimensions, we discretize the motor speeds into 25 bins. For each bin, we compute $p_{profile}(v_1, v_2) = \max_{v_3} p(v_1, v_2, v_3)$, where $p(v_1, v_2, v_3)$ is the performance of the robot for motor speeds v_1, v_2, v_3 and $p_{profile}(v_1, v_2)$ is the performance reported in the profile. To get a comprehensive pictures, we need three plots: $p_{profile}(v_1, v_2)$, $p_{profile}(v_1, v_3)$, and $p_{profile}(v_2, v_3)$.

Learning algorithm

Our learning algorithm allows the robot to discover by trial-and-error the best rotation speeds for its three motors. It essentially implements a variant of Bayesian optimization, which is a state-of-the-art optimization algorithm designed to maximize expensive performance functions (a.k.a. cost functions) whose gradient cannot be evaluated analytically[27,28]. Like other model-based optimization algorithms (e.g., surrogate-based algorithms[37–39], kriging[40], or DACE[41,42]), Bayesian optimization models the objective function with a regression method, uses this model to select the next point to acquire, then updates the model, etc. until the algorithm has exhausted its budget of function evaluations.

Here a Gaussian process models the objective function[29], which is a common choice for Bayesian optimization[27,28,43–45]. For an unknown cost function f, a Gaussian process defines the probability distribution of the possible values $f(\mathbf{x})$ for each point $\mathbf{x}$. These probability distributions are Gaussian, and are therefore defined by a mean (μ) and a variance (σ^2). However, μ and σ^2 can be different for each $\mathbf{x}$; a Gaussian process therefore defines a probability distribution *over functions*:

$$P(f(\mathbf{x})|\mathbf{x}) = \mathcal{N}(\mu(\mathbf{x}), \sigma^2(\mathbf{x})) \tag{1}$$

where $\mathcal{N}$ denotes the standard normal distribution.

At iteration t, if the performance $[P_1, \cdots, P_t] = \mathbf{P}_{1:t}$ of the points $[\chi_1, \cdots, \chi_\mathbf{t}] = \chi_{1:t}$ has already been evaluated, then $\mu_t(\mathbf{x})$ and $\sigma_t^2(\mathbf{x})$ are fitted as follows[29]:

$$\begin{aligned}
&\mu_t(\mathbf{x}) = \mathbf{k}^\intercal \mathbf{K}^{-1} \mathbf{P}_{1:t} \\
&\sigma_t^2(\mathbf{x}) = k(\mathbf{x}, \mathbf{x}) + \sigma_{noise}^2 - \mathbf{k}^\intercal \mathbf{K}^{-1} \mathbf{k} \\
&\text{where:} \\
&\mathbf{K} = \begin{bmatrix} k(\chi_1, \chi_1) & \cdots & k(\chi_1, \chi_t) \\ \vdots & \ddots & \vdots \\ k(\chi_t, \chi_1) & \cdots & k(\chi_t, \chi_t) \end{bmatrix} + \sigma_{noise}^2 I \\
&\mathbf{k} = \begin{bmatrix} k(\mathbf{x}, \chi_1) & k(\mathbf{x}, \chi_2) & \cdots & k(\mathbf{x}, \chi_t) \end{bmatrix}
\end{aligned} \tag{2}$$

The matrix $\mathbf{K}$ is called the covariance matrix. It is based on a *kernel function* $k(\mathbf{x_1}, \mathbf{x_2})$ which defines how samples influence each other. Kernel functions are classically variants of the Euclidean distance. Here we use the *exponential kernel*[15,28,29,44]:

$$k(\mathbf{x_1}, \mathbf{x_2}) = \exp\Big(-\frac{1}{\beta^2}||\mathbf{x_1} - \mathbf{x_2}||^2\Big) \quad (3)$$

We fixed β to 0.15.

An interesting feature of Gaussian processes is that they can easily incorporate a prior $\mu_p(\mathbf{x})$ for the mean function, which helps to guide the optimization process to zones that are known to be promising:

$$\mu_t(\mathbf{x}) = \mu_p(\mathbf{x}) + \mathbf{k}^\intercal \mathbf{K}^{-1}(\mathbf{P}_{1:t} - \mu_p(\chi_{1:t})) \quad (4)$$

In our implementation, the prior is a second Gaussian process defined by hand-picked points (see the "prior" section below).

To select the next χ to test (χ_{t+1}), Bayesian optimization maximizes an *acquisition function*, a function that reflects the need to balance exploration – improving the model in the less known parts of the search space – and exploitation – favoring parts that the model predicts as promising. Numerous acquisition functions have been proposed (e.g., probability of improvement, the expected improvement, or the Upper Confidence Bound (UCB)[28,43,44]); we chose UCB because it provided the best results in several previous studies[43,44] and because of its simplicity. The equation for UCB is:

$$\chi_{t+1} = \arg\max_{\mathbf{x}}(\mu_t(\mathbf{x}) + \kappa\sigma_t(\mathbf{x})) \quad (5)$$

where κ is a user-defined parameter that tunes the tradeoff between exploration and exploitation. We chose $\kappa = 0.2$.

Prior for the learning algorithm

The learning algorithm is guided by a prior that captures the idea that the highest-performing gaits are likely to be a combination of motors at full speed (in forward or in reverse). In our implementation, it is implemented with a Gaussian process defined by 9 hand-picked points and whose variance is ignored (equation 2). The kernel function is the exponential kernel (equation 3), with $\beta = 0.15$.

The 9 hand-picked points ($\chi_1, \cdots, \chi_9$) are as follows (Fig 5-D):

$\chi_1 = [-100\%, -100\%, -100\%]$ $\quad$ $\chi_2 = [-100\%, -100\%, +100\%]$
$\chi_3 = [-100\%, 100\%, -100\%]$ $\quad$ $\chi_4 = [-100\%, +100\%, +100\%]$
$\chi_5 = [+100\%, -100\%, -100\%]$ $\quad$ $\chi_6 = [+100\%, +100\%, +100\%]$
$\chi_7 = [+100\%, -100\%, -100\%]$ $\quad$ $\chi_8 = [+100\%, -100\%, +100\%]$
$\chi_9 = [0\%, 0\%, 0\%]$;
$P(\chi_1), \cdots, P(\chi_8) = 0.3; P(\chi_9) = 0$

Statistics

For all experiments, we report the 5^{th} and 95^{th} percentiles. We used a two-tailed Mann-Whitney U test for all statistical tests. For the box plots, the central mark is the median, the edges of the box are the 25$^{\text{th}}$ and 75$^{\text{th}}$ percentiles, the whiskers extend to the most extreme data points not considered outliers, and outliers are plotted individually. For each box plot, the result of the Mann-Whitney U test (two-tailed) is indicated with stars: * means $p \leq 0.05$, ** means $p \leq 0.01$, *** means $p \leq 0.001$, and **** means $p \leq 0.0001$.

Computer code

`http://members.loria.fr/JBMouret/src/limbo-tensegrity.tar.gz` ; this code will be released with an open-source license on Github for the final publication.

Data availability

`http://members.loria.fr/JBMouret/data/tensegrity.tar.gz` ; these data will be released on Dryad for the final publication.

1. Sumbre G, Fiorito G, Flash T, Hochner B (2005) Neurobiology: motor control of flexible octopus arms. *Nature* 433(7026):595–596.
2. Gemmell BJ et al. (2013) Passive energy recapture in jellyfish contributes to propulsive advantage over other metazoans. *Proceedings of the National Academy of Sciences* 110(44):17904–17909.
3. Simon MA et al. (2010) Visceral-locomotory pistoning in crawling caterpillars. *Current biology* 20(16):1458–1463.
4. Snelson K (2012) The art of tensegrity. *International Journal of Space Structures* 27(2-3):71–80.
5. Skelton RE, de Oliveira MC (2009) *Tensegrity systems*. (Springer) Vol. 1.
6. Wang N et al. (2001) Mechanical behavior in living cells consistent with the tensegrity model. *Proceedings of the National Academy of Sciences* 98(14):7765–7770.
7. Lipson H (2014) Challenges and opportunities for design, simulation, and fabrication of soft robots. *Soft Robotics* 1(1):21–27.
8. Wehner M et al. (2016) An integrated design and fabrication strategy for entirely soft, autonomous robots. *Nature* 536(7617):451–455.
9. Trimmer B (2008). *Applied Bionics and Biomechanics* 5(3).
10. Lin HT, Leisk G, Trimmer B (2011) Goqbot: A caterpillar-inspired soft-bodied rolling robot. *Bioinspiration and Biomimetics* 6:026007.
11. Shepherd RF et al. (2011) Multigait soft robot. *Proceedings of the National Academy of Sciences* 108(51):20400–20403.
12. Martinez RV et al. (2013) Robotic tentacles with three-dimensional mobility based on flexible elastomers. *Advanced Materials* 25(2):205–212.
13. Bartlett NW et al. (2015) A 3d-printed, functionally graded soft robot powered by combustion. *Science* 349(6244):161–165.
14. Craig J (1989) *Introduction to Robotics*. (Addison-Wesley).
15. Cully A, Clune J, Tarapore D, Mouret JB (2015) Robots that can adapt like animals. *Nature* 521(7553):503–507.
16. Valero-Cuevas FJ et al. (2007) The tendon network of the fingers performs anatomical computation at a macroscopic scale. *IEEE Transactions on Biomedical Engineering* 54(6):1161–1166.
17. Ceceri K (2015) *Making Simple Robots: Exploring Cutting-Edge Robotics with Everyday Stuff*. (Maker Media, Inc.).
18. Caluwaerts K et al. (2014) Design and control of compliant tensegrity robots through simulation and hardware validation. *Journal of The Royal Society Interface* 11(98):20140520.
19. Koizumi Y, Shibata M, Hirai S (2012) Rolling tensegrity driven by pneumatic soft actuators in *Proc. of the IEEE Int. Conf. on Robotics and Automation (ICRA)*. (IEEE), pp. 1988–1993.
20. Bliss TK, Iwasaki T, Bart-Smith H (2012) Resonance entrainment of tensegrity structures via cpg control. *Automatica* 48(11):2791–2800.
21. Mirletz BT et al. (2015) Goal-directed cpg-based control for tensegrity spines with many degrees of freedom traversing irregular terrain. *Soft Robotics* 2(4):165–176.
22. Jakobi N, Husbands P, Harvey I (1995) Noise and the reality gap: The use of simulation in evolutionary robotics in *European Conference on Artificial Life*. (Springer), pp. 704–720.
23. Koos S, Mouret JB, Doncieux S (2013) The transferability approach: Crossing the reality gap in evolutionary robotics. *IEEE Transactions on Evolutionary Computation* 17(1):122–145.
24. Oppenheim IJ, Williams WO (2001) Vibration of an elastic tensegrity structure. *European Journal of Mechanics-A/Solids* 20(6):1023–1031.
25. Khazanov M, Humphreys B, Keat W, Rieffel J (2013) Exploiting dynamical complexity in a physical tensegrity robot to achieve locomotion. *Advances in Artificial Life, ECAL* 12:965–972.
26. Khazanov M, Jocque J, Rieffel J (2014) Evolution of locomotion on a physical tensegrity robot in *ALIFE 14: The Fourteenth Conference on the Synthesis and Simulation of Living Systems*. pp. 232–238.
27. Ghahramani Z (2015) Probabilistic machine learning and artificial intelligence. *Nature* 521(7553):452–459.
28. Shahriari B, Swersky K, Wang Z, Adams RP, de Freitas N (2016) Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE* 104(1):148–175.
29. Rasmussen CE, Williams CKI (2006) *Gaussian processes for machine learning*. (MIT Press).
30. Davison AJ, Reid ID, Molton ND, Stasse O (2007) Monoslam: Real-time single camera slam. *IEEE transactions on pattern analysis and machine intelligence* 29(6):1052–1067.

31. Mouret JB, Clune J (2015) Illuminating search spaces by mapping elites. *arXiv:1504.04909 [cs, q-bio]*. arXiv: 1504.04909.
32. Reuillon R, Schmitt C, De Aldama R, Mouret JB (2015) A new method to evaluate simulation models: The calibration profile (cp) algorithm. *Journal of Artificial Societies and Social Simulation* 18(1):12.
33. Rubenstein M, Ahler C, Nagpal R (2012) Kilobot: A low cost scalable robot system for collective behaviors in *Proc. of the IEEE International Conference Robotics and Automation (ICRA)*. (IEEE), pp. 3293–3298.
34. Romanishin JW, Gilpin K, Rus D (2013) M-blocks: Momentum-driven, magnetic modular robots in *Proc. of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. (IEEE), pp. 4288–4295.
35. Chase R, Pandya A (2012) A review of active mechanical driving principles of spherical robots. *Robotics* 1(1):3–23.
36. Pfeifer R, Lungarella M, Iida F (2007) Self-organization, embodiment, and biologically inspired robotics. *science* 318(5853):1088–1093.
37. Booker AJ et al. (1999) A rigorous framework for optimization of expensive functions by surrogates. *Structural optimization* 17(1):1–13.
38. Forrester AIJ, Keane AJ (2009) Recent advances in surrogate-based optimization. *Progress in Aerospace Sciences* 45(1):50–79.
39. Jin Y (2011) Surrogate-assisted evolutionary computation: Recent advances and future challenges. *Swarm and Evolutionary Computation* 1(2):61–70.
40. Simpson TW, Mauery TM, Korte JJ, Mistree F (1998) Comparison of response surface and kriging models for multidisciplinary design optimization. *American Institute of Aeronautics and Astronautics* 98(7):1–16.
41. Jones DR, Schonlau M, Welch WJ (1998) Efficient global optimization of expensive black-box functions. *Journal of Global optimization* 13(4):455–492.
42. Sacks J, Welch WJ, Mitchell TJ, Wynn HP, et al. (1989) Design and analysis of computer experiments. *Statistical science* 4(4):409–423.
43. Calandra R, Seyfarth A, Peters J, Deisenroth MP (2014) An experimental comparison of bayesian optimization for bipedal locomotion in *Proc. of the IEEE International Conference on Robotics and Automation (ICRA)*.
44. Brochu E, Cora VM, De Freitas N (2010) A tutorial on bayesian optimization of expensive cost functions, with application to active user modeling and hierarchical reinforcement learning. *arXiv preprint arXiv:1012.2599*.
45. Lizotte DJ, Wang T, Bowling MH, Schuurmans D (2007) Automatic gait optimization with Gaussian process regression. in *Proceedings of the the International Joint Conference on Artificial Intelligence (IJCAI)*. Vol. 7, pp. 944–949.

Acknowledgments

This project received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Project: ResiBots, grant agreement No 637972); J.R. received funding for a 1-month visit at University of Lorraine, France. The authors would like to thank Dorian Goepp for his invaluable help and the ResiBots team for their comments on this manuscript.

Author Contributions

J.R. designed the robot; J.R. and J.-B.M. designed the study, performed the experiments, analyzed the data, and wrote the paper.

Supplementary information

Video S1

Presentation of our soft tensegrity robot. The video shows the soft tensegrity robot in action: how it can locomote and how it can learn to compensate when damaged. This video is available on Youtube: `https://www.youtube.com/watch?v=SuLQDhrk9tQ`